

**MAINTAINING SESSION CONNECTIVITY WHEN A MOBILE NODE
MOVES FROM ONE LAYER 3 NETWORK TO ANOTHER**

Inventor

Aseem SETHI

No. 8, Raymond Court, 23 MSO Colony
Jeevanhalli Main Road
Cox Town, Bangalore
Karnataka (State), India - 560 005
Citizenship: India

Assignee:

Cisco Technology, Inc.
A California Corporation
170 W. Tasman Drive
San Jose, CA 95134
Telephone: (408) 525-9706
Fax: (408) 526-5952

Attorney:

Law Firm of Naren Thappeta
Phone: +1 (510) 342-2519 x-6580
Fax: + 1 (707) 356-4172
Email: naren@iphorizons.com
URL: www.iphorizons.com

MAINTAINING SESSION CONNECTIVITY WHEN A MOBILE NODE MOVES FROM ONE LAYER 3 NETWORK TO ANOTHER

Background of the Invention

Field of the Invention

5 The present invention relates to mobile communication applications, and more specifically to a method and apparatus for maintaining session connectivity when a mobile node moves from one layer 3 network to another.

Related Art

10 Mobile nodes refer to devices such as cell phones which are often carried by users from one physical location to another. Mobile nodes are used to communicate with other devices (generally referred to as correspondent nodes) using wireless protocols as is well known in the relevant arts. The communication often forms the basis for applications such as internet multimedia conferences, internet telephone calls and multimedia distribution.

15 Sessions are generally maintained to support applications between mobile nodes and correspondent nodes as is well known in the relevant arts. Session Initiation Protocol (SIP) is a common application-layer control (signaling) protocol for creating, modifying and terminating sessions as is also well known in the relevant arts. SIP is described in further detail in RFC 2543 entitled, "SIP: Session Initiation Protocol", available from www.ietf.org, and is incorporated in its entirety herewith.

20 Sessions in turn are often supported by layer-3 (e.g., Internet Protocol (IP)) network protocols. As is also well known, each mobile node needs to be assigned a layer-3 address. The layer-3 address

is typically used by the corresponding correspondent node to send layer-3 packets to the mobile node. The network address portion of layer-3 network is used to route the packets to the mobile node.

Mobile nodes may be moved from one physical location to another. In general, it is desirable to provide seamless and uninterrupted support to applications executing on the mobile nodes even when mobile nodes move from one physical location to the other.

One feature that poses challenges to such support is the fact that devices in different geographical locations often operate using different layer-3 network address spaces for the devices. In general, a node in a layer-3 network can send packets only if the node is assigned a layer-3 address in the address space corresponding to the layer-3 network.

Protocols such as SIP provide a mechanism for a mobile node to operate in different locations served by different layer-3 address spaces. Typically, a mobile node operating at a home location registers with a SIP server ("home server") of that location. The information on the home server can then be used by correspondent nodes to initiate communication with the mobile node. Irrespective of which node initiates communication, a session is established with a correspondent node using the layer-3 address assigned to the mobile node in the location.

The mobile node can be moved to a new location served by a different layer-3 address space. The mobile node is assigned a new layer-3 address in the different layer-3 address space. The mobile node registers with a SIP proxy server in the new location. Registration generally entails providing an indication of the move and typically also notifying the SIP proxy server of the new IP address.

The SIP proxy server may communicate with the home proxy server to notify the home proxy server of the presence of the mobile node in the new location. Any necessary information may be exchanged between the two servers to enable correspondent nodes to communicate with the mobile node in the new location. Thus, the correspondent nodes can use the information in the home server and/or the proxy server to initiate applications afresh with the mobile node in the new location. The new layer-3 address may be used for the mobile node in the new location.

Thus, using protocols such as SIP, a mobile node may operate from different locations having different IP addresses. However, the approach(es) of above would operate accurately only when sessions (or applications) are started afresh in the new location, but not for sessions ("active session") which may be active when a mobile node moves to a new location. In other words, for sessions initiated after the move, the new layer-3 address would be used consistently by both the mobile node and the correspondent node, and communication would thus established accurately.

On the other hand, for sessions which are active when a mobile node moves to the new location, one or both of the mobile node and correspondent node may be designed to continue to operate with the old address of the mobile node. The network (between the two nodes) may not be able to accurately deliver packets to the mobile node as the packets may be routed to the old layer-3 network.

As a result, the packets transmitted by the correspondent node may not be received by the mobile node. Consequently, sessions (and thus applications) which are active prior to a move may fail (terminate) when a mobile node moves from one geographical location to another, which is often

undesirable. Therefore, what is needed is a method and apparatus for maintaining session connectivity when a mobile node moves from one layer-3 network to another.

Summary of the Invention

5 An aspect of the present invention enables a session between a mobile node and a correspondent node to continue without terminating even if the mobile node is assigned a new address upon moving to a new location. Such a feature is achieved by performing network address translation (NAT) between an old address (assigned in a old location prior to the move) and the new address in network devices in a path between the mobile node and the correspondent node.

10 For example, a first hop router in the path from the correspondent node to mobile node is configured to replace a old address in the destination address field of a packet (received from mobile node) with a corresponding new address. As a result, the correspondent node can continue to send packets to mobile node with the old address, as before the move. Due to the use of the new address, the packet is routed accurately to the new location.

15 In an embodiment, another router at the new location may be configured to replace the new address with the old address in the destination address field. A static route may also be configured on the another router to ensure that the packet with the old address is delivered to the mobile node (even though the mobile node has been assigned the new address). Thus, packets are sent (at correspondent node) and delivered (to mobile node) using the old address as before the move, enabling the sessions to continue without terminating.

According to another aspect of the present invention, a mobile node sends both the old address and the new address in a registration message to a proxy server in the new location. By using the two addresses, the proxy server may configure the router at the first hop from the mobile node.

The mobile node may further determine a correspondent node related to each active session, and send the old and new addresses to a mobile server (e.g., a Session Initiation Protocol (SIP) server) serving the correspondent node. The server in turn may configure the first hop router in the path from the correspondent node to mobile node. Once the two routers are configured, the packets related to sessions active before the move, may be sent and delivered using the old address.

Further features and advantages of the invention, as well as the structure and operation of various embodiments of the invention, are described in detail below with reference to the accompanying drawings. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements. The drawing in which an element first appears is indicated by the leftmost digit(s) in the corresponding reference number.

Brief Description of the Drawings

The present invention will be described with reference to the accompanying drawings, wherein:
Figure 1 is a block diagram illustrating the manner in which basic terminal mobility can be achieved using SIP;

Figure 2 is a flow chart illustrating a method which enables a mobile node and a correspondent node to continue to operate with the old IP address of the mobile node even after the mobile node moves;

Figure 3A is a flow chart illustrating a method using which a SIP server can be implemented to configure network devices to support terminal mobility according to an aspect of the present invention;

Figure 3B is a flow chart illustrating a method using which a mobile node initiates configurations of various network devices to enable terminal mobility according to an aspect of the present invention;

Figure 4 is a block diagram illustrating an embodiment of the present invention implemented substantially in the form of software; and

Figure 5 is a block diagram illustrating an embodiment of server implemented substantially in the form of hardware.

Detailed Description of the Preferred Embodiments

1. Overview and Discussion of the Invention

An aspect of the present invention allows an active session between a mobile node and a correspondent node to continue without failing (terminating) even if the mobile node is assigned a new layer-3 address when the mobile node moves from a home location to a new location. The feature may be achieved by translating the new layer-3 address (in the new location) to a old layer-3 address (in the home location), and vice versa.

As a result, a correspondent node can continue to operate using the same old layer-3 address even after the mobile node moves to the new location, and active sessions terminating on the mobile node can continue to be supported. By appropriate implementation of other components, the mobile node can also be made to operate such that any sessions with the correspondent node also do not fail.

It should be understood that numerous specific details, relationships, and methods are set forth to provide a full understanding of the invention. One skilled in the relevant art, however, will readily recognize that the invention can be practiced without one or more of the specific details, or with other methods, etc. In other instances, well-known structures or operations are not shown in detail to avoid obscuring the invention.

Several aspects of the invention are described below with reference to example environments for illustration. It is useful to understand the basic terminal mobility provided by SIP to appreciate several aspects of the present invention. Accordingly, basic terminal mobility is described first.

2. Basic Terminal Mobility in an Example Environment

Figure 1 is a block diagram illustrating the manner in which session initiation protocol (SIP) generally supports basic terminal mobility (i.e., for a mobile node to operate from different locations). Figure 1 is shown containing networks 111, 151 and 191, mobile node 110 (shown as 110-A in location 101 and 110-B in location 102), routers 120-A through 120-C, servers 130, 140 and 170, internet 150, and correspondent node 160. Each component is described in further detail below.

Mobile node 110-A, sub-network 111, router 120-A and server 130 are shown in location 101. Mobile node 110-B (same as mobile node 110-A, but in location 102), sub-network 151, router 120-B and server 140 are shown in location 102. Locations 101 and 102 are respectively referred to as old location and new location as the same mobile node is described as being moved from location 101 to location 102. Correspondent node 160, router 120-C, server 170 and sub-network 191 are shown in location 103 ("correspondent location").

Networks 111, 151 and 191 represent layer-3 networks in locations 101, 102 and 103 respectively. Each network (e.g., 111) provides connectivity among the connected components (e.g., server 130, router 120-A and mobile node 110). Routers 120-A, 120-B and 120-C are respectively shown connecting networks 111, 151 and 191 to internet 150. Networks 111, 151 and 191, routers 120-A, 120-B and 120-C and Internet 150 may be implemented using Internet Protocol (IP) in a known way.

Mobile node 110-A (in old location 101) represents a device such as a cell phone which is often carried by a user from one location to the other. Mobile node 110-A may communicate with correspondent node 160 and server 130 using a layer-3 protocol. Mobile node 110-A is shown accessing network 151 by using a wireless protocol (shown with a broken arrow). Mobile node 110-A is assigned a layer-3 address ("old address") in an address space assigned to network 111. The address may be assigned in a known way.

Server 130 represents a mobile server which generally assists mobile nodes to have terminal mobility. In an embodiment, server 130 is implemented consistent with SIP protocol (RFC 2543). For illustration, it is assumed that server 130 operates as a 'home server' (also referred to as S-CSCF (serving call session control function)) for mobile node 110-A. In general, all mobile nodes (including 110-A) in location 101 may be registered with server 130 operating as a home server. When mobile node 110-A moves to new location 102, server 130 redirects SIP requests related to mobile node 110-A to server 140 to support basic terminal mobility.

The manner in which basic terminal mobility may be supported is described below with a

situation in which mobile node 110-A moves to new location 102. The same mobile node is referred to as mobile node 110-B for clarity.

Mobile node 110-B is assigned a new layer-3 address in an address space corresponding to new location 102. The assignment may be performed in a known way using protocols such as dynamic host configuration protocol (DHCP). Mobile node 110-B registers the new address with proxy server 140.

Server 140 operates as a proxy server (P-CSCF, proxy call session control function) for mobile node 101-B in new location 102. Server 140 receives registration information including the old address and new address from mobile node 101-B. Server 140 determines the home server for mobile node 101-B (for example, by using DNS query, as is well known in the relevant arts) and communicates the new address for mobile node 110-A/B to home server 130.

Any new correspondent nodes can then initiate sessions with mobile node 110-B (in new location 102). To initiate a session/application with mobile node 110-B, correspondent node 160 generally determines the home proxy server for mobile node 110-B by sending an identifier (Uniform Resource Identifier in SIP environment) to server 170. Server 170 sends the identifier to a domain name system (DNS, not shown) to determine the home server of mobile node 110-B. DNS returns a layer-3 address of the home server of mobile node 110-B as server 130.

Server 170 sends a SIP request to server 130 to determine where mobile node 110-B is presently located. Server 130 redirects the SIP request to server 140 (based on the information

exchanged after mobile node 110-A has moved to new location 102) and server 140 sends a SIP response to server 170 providing a new layer-3 address for mobile node 110-B. Correspondent node 160 may now initiate a session with mobile node 110-B using the provided layer-3 address.

Thus, correspondent node 160 can initiate sessions/applications to mobile node 110-B after mobile node 110-B has moved to new location 102. Similarly, mobile node 110-B can initiate sessions/applications to correspondent node 160 using the new address assigned at new location 102. Thus, mobile node 110-B may be said to have basic terminal mobility.

The approach described above would operate accurately when session are started afresh between mobile node 110-B and correspondent node 160. However, the above approach may not work for sessions ("active sessions") which are active/alive when mobile nodes move from old location 101 to new location 102. In particular, the packets transmitted by correspondent node 160 may not be received by mobile node 110-B causing any active sessions to terminate.

Accordingly, an approach using which active sessions may continue to operate without termination is described below.

3. Enabling Active Sessions to Continue After Mobile Node Moves

Figure 2 is a flow-chart illustrating a method by which any active sessions may continue without termination when a mobile node moves to a new location. The method is described below with reference to Figure 1 for illustration. However, the method may be implemented in other environments as well. The method begins in step 201, in which control immediately passes to step 210.

In step 210, the active sessions terminating on mobile node 110-B are determined. In an embodiment implemented in the context of Internet Protocol, the active sessions are determined by examining tables (e.g., active TCP/RCP ports) stored internally in mobile node 110-B. The determination can be performed in a known way. As may be appreciated, each active session has mobile node 110-B at one end and a correspondent node at another end. The rest of the steps are described with reference to correspondent node 160 as being at the other end of a session.

In step 220, a network device on the path from each correspondent node to mobile node 110-B is configured to translate old address to new address in the destination address field of each layer-3 packet. In an embodiment, the network device corresponds to first hop router 120-C present in the path from correspondent node 160 to mobile node 110-B.

By configuring router 120-C, correspondent node 170 can continue to send layer-3 packets with the old address (of location 101) in the destination field as the old address is immediately replaced with the new address in router 120-C at the first hop. The packets containing the new address may then be accurately routed to network 151 now containing mobile node 110-B. As a result, correspondent node 160 may continue to operate with the same view of the sessions both before and after mobile node 110-A/B moves, and the session is not terminated by correspondent node 160.

Another aspect of the present invention enables mobile node 110-B to receive packets with the old layer-3 address in the destination address field. Step 230 may be implemented to enable mobile node 110-B to receive packets with the old layer-3 address.

In step 230, another network device is configured to translate the new address back to old address in each layer-3 packet before delivery to mobile node 110-B. In an embodiment, network device corresponds to router 120-B which is also the closest (last hop) router to mobile node 110-B in location 102. One problem with such a configuration is that packets with old address in the destination field may be routed to old location 101. Accordingly, step 240 may be performed to ensure that the packets are delivered to mobile node 110-B in new location 102.

In step 240, network devices are configured to deliver packets with the old address (in destination field) to mobile node 110-B in new location 102. When router 120-B alone is configured in step 230, a static route may be configured in router 120-B to cause the packets with destination address of old address to be delivered to mobile node 110-B on network 151. The static route may be configured using protocols such as SNMP in a know way.

As noted above, network devices (routers 120-B and 120-C) may need to be configured to enable the sessions to continue to operate without termination after mobile node 101-A moves. The manner in which the configurations can be initiated/caused is described below.

4. Configuring Network Devices

Figure 3A is a flow chart illustrating a method by which the first hop routers 120-B and 120-C may be configured to perform network address translation. The method is described below first with reference to server 140 of Figure 1. The manner in which the method can be implemented in server 170 is described later. The method begins in step 301, in which control immediately passes to step 310.

In step 310, server 140 receives a message containing the old and new addresses of mobile node 110-A. The message may be received as a part of SIP registration, which is usually performed when a mobile node moves to a new location as described in further detail in RFC 2543. In an embodiment, the registration message format is modified to include the old address as described in a section below.

In step 320, server 140 stores a mapping entry in network address translation (NAT) table of router 120-B to cause router 120-B to translate the new address to the old address in the destination field of the packets received from Internet 150. The configuration may be performed using protocols such as SNMP in a known way. Server 140 may also configure a static route within router 120-B to send packets with the old address on network 151 to mobile node 110-B in new location 102.

Similarly, server 170 may also receive a message containing old and new addresses of mobile node 110-A, and configure a NAT table in router 120-C to cause the old address to be translated to new address in the destination address field of packets received from correspondent node 160. In an embodiment, each of the routers 120-A, 120-B and 120-C are respectively co-located (or integrated into one box) with a corresponding one of the servers 130, 140 and 170. Due to the co-location, reconfiguration of routers merely entails internal reconfiguration (typically of different tables) of the integrated boxes.

As described above, mobile node 110-B initiates configuration of routers 120-B and 120-C by sending appropriate messages to corresponding proxy/home servers. The operation of mobile node 110-B is summarized below with reference to the flow chart of Figure 3B.

5. Mobile Node

Figure 3B is a flow-chart illustrating a method using which a mobile node can cause configuration of various network devices. The method begins in step 351, in which control passes to step 360. In step 360, mobile node 110-B determines that it (mobile node 110-B) has been moved to new location 102. The determination can be performed in one of several known ways. Some examples are noted in the below paragraph.

In embodiments implemented consistent with SIP, mobile node 110-B may be implemented to determine presence in a new network by examining network notifications (e.g., ICMP message from foreign router 120-B). Alternatively, mobile node 110-B may periodically send SIP register requests and determine presence in a new network upon receiving a register reply message from a new server (server 140 instead of server 130).

In step 370, mobile node 110-B initiates assignment of a new IP address consistent with a protocol implemented in new location 102. Applications such as dynamic host configuration protocol (DHCP) may be used for such an assignment. When assigned, the new IP address is received by mobile node 110-B from a DHCP server (not shown).

In step 380, mobile node 110-B sends a registration message to proxy server 140 at new location 102. The proxy server may be determined in a known way and the new address may be used to send the registration message. The old address (used in old location 101) may also be sent along with the registration as described below with reference to an example embodiment(s).

In step 390, mobile node 110-B sends a notification message indicating the move and the two (old and new) addresses to proxy (or home) server 170 of correspondent node 160. Server 170 then causes configuration of router 120-C.

Once the configurations of routers 120-B and 120-C is complete, the nodes at either end of a session may continue to transmit packets as if move did not occur as described below with reference to an example.

6. Example

The approaches of Figures 2, 3A and 3B are illustrated with an example now. For illustration, it is assumed that old layer-3 (IP) address of mobile node 110-A (in location 101) is 1.1.1.1, and the address of the correspondent node 160 (in location 103) is 3.3.3.3.

Packets sent mobile node 110-A to correspondent node 160 will generally have 1.1.1.1 as the source address and 3.3.3.3 as the destination address. On the other hand, packets sent from correspondent node 160 to mobile node 110-A will have a source address of 3.3.3.3 and destination address of 1.1.1.1. As described below, the mobile node 101-A and correspondent node 160 may continue to send and receive packets with the same IP addresses even after mobile node 110-A moves, as described below.

Mobile node 110-A/B is assumed to be assigned a new layer-3 address 2.2.2.2 when moved to new location 102. Mobile node 110-B registers with proxy server 140 present in new location 102 by sending a register request. The old address 1.1.1.1 and the new address 2.2.2.2 may also be

passed to server 140.

Server 140 configures router 120-B to translate the old address to new address in the destination field of packets received from Internet 150. Similarly, server 170 also receives the old address and new address, and configures router 120-C to translate the old address (1.1.1.1) to the new address (2.2.2.2). A static route to send packets with the destination address of the old address to network 151 may also be configured in router 120-B as noted above with reference to step 240 of Figure 2.

Once the configurations are complete, packets sent by correspondent node 160 with the old address (1.1.1.1) in the destination field may be delivered to mobile node 110-B as described now. Router 120-C performs network address translation to replace the old address 1.1.1.1 with new address 2.2.2.2. The new address causes the packet to be routed to router 120-B. Router 120-B performs network address translation again to replace the new address 2.2.2.2 with old address 1.1.1.1 and routes the packet to mobile node 110-B due to the static route.

Thus, correspondent node 160 may continue sending packets to mobile node 110-B with the old address 1.1.1.1 as destination address, even though mobile node 110-B has been assigned a new layer-3 address of 2.2.2.2 in new location 102. Similarly, mobile node 110-B may continue to send packets with a source address of 1.1.1.1, and the packets would get routed to correspondent node 160 as the source address is generally not examined during routing decisions.

While the embodiments of above are described as performing NAT only for packets sent from

correspondent node 160 to mobile node 110-B, the first hop routers 120-B and 120-C may be configured to perform NAT in the other direction as well. In such a scenario, router 120-B replaces the old IP address (1.1.1.1) in the source address field with the new IP address (2.2.2.2), and the packet with the new address is transmitted via Internet 150 to router 120-C. Router 120-C replaces the new IP address with the old IP address before forwarding to correspondent node 160.

The description is continued with reference to the format of the packet when mobile node 110-B registers with proxy server 140.

7. Packet Format

When mobile node 110-B moves to new location 102, mobile node 110-B registers with server 140 by sending a register request. When sending the register request, mobile node 110-B passes the old address corresponding to old location 101 to server 140. In an embodiment implemented consistent with RFC 2543 (which is incorporated in its entirety herewith), an optional tag may be used to communicate the old address of mobile node 110-B. Assuming that old and new addresses of mobile node 110-A/B are 1.1.1.1 and 2.2.2.2 respectively, an example register message may contain the following information:

```
REGISTER sip:p-cscf.cisco.com SIP/2.0
Require: com.cisco.TerminalMobilityUsingNat
Via: SIP/2.0/UDP p-cscf.com:5060
From: MN <sip:MN@ p-cscf.com>
To: MN <sip:MN@ p-cscf.com>
Call-ID: 123456789@ p-cscf.com
CSeq: 1 REGISTER
Contact: <sip:MN@2.2.2.2>
Content-Length: 0
com.cisco.TerminalMobilityUsingNat: <sip:MN@1.1.1.1>
```

Only the lines which communicate the old address are described here for conciseness. The other lines can be interpreted consistent with RFC 2543. The Require statement (in line 2 above) indicates that a new optional tag with a URL of com.cisco.TerminalMobilityUsingNat is being used. The last statement provides a value of 1.1.1.1 as the old IP address (for the option tag referenced in the Require statement).

Similarly, mobile node 110-B may notify server 170 (the proxy or home server of correspondent node 160) of the move to new location and the new and old IP addresses. Any approaches may be used for such notifications. At least in situations when mobile node 110-A has initiated the active session using a SIP_invite packet (described in RFC 2543), the packet may be re-sent with the old address also being included in the invite packet. The same optional tag and format described above may be used to communicate the old address.

In other situations, mobile node 110-B may determine the presently serving home/proxy server of correspondent node 160 in a known way, and send the old address to server 170. The old address may be sent in a notification message which further in effect requests configuration of router 120-C at the first hop in the path from correspondent node 160 to mobile node 110-B.

Once the old IP address is communicated to SIP servers 140 and 170, the first hop routers may be configured as described above. The sessions which are active before mobile node 110-A is moved to new location 102, may then continue to operate without termination.

It should be understood that each feature of the present invention can be implemented in a combination of one or more of hardware, software and firmware. In general, when throughput performance is of primary consideration, the implementation is performed more in hardware (e.g., in the form of an application specific integrated circuit).

5 When cost is of primary consideration, the implementation is performed more in software (e.g., using a processor executing instructions provided in software/firmware). Cost and performance can be balanced by implementing routers 120-A through 120-C with a desired mix of hardware, software and/or firmware. An embodiment of router 120-A implemented substantially in software is described below.

10 8. Software Implementation

Figure 4 is a block diagram illustrating the details of system 400, which may represent one of servers 140 and 170, and mobile node 110-B. System 400 is shown containing processing unit 410, random access memory (RAM) 420, storage 430, output interface 460, network interface 480 and input interface 490. Each component is described in further detail below.

15 Output interface 460 provides output signals (e.g., display signals to a display unit, not shown) which can form the basis for a suitable user interface for an administrator to interact with system 400. Input interface 490 (e.g., interface with a key-board and/or mouse, not shown) enables an administrator to provide any necessary inputs to system 400.

Output interface 460 and input interface 490 can be used, for example, to enable a network

administrator to configure system 400 to provide various features of the present invention. For example, when system 400 corresponds to mobile node 110-B, an administrator may either enable or disable the feature of allowing the sessions to continue to operate upon move to a new location. When system 400 corresponds to one of servers 140 and 170, the administrator may specify the specific mechanism (e.g., SNMP) to be used in configuring the corresponding routers.

Network interface 480 enables system 400 to send and receive data on communication networks using internet protocol (IP). Network interface 480, output interface 460 and input interface 490 can be implemented in a known way.

RAM 420 and storage (secondary memory) 430 may together be referred to as a memory. RAM 420 receives instructions and data on path 450 from storage 430, and provides the instructions to processing unit 410 for execution. In addition, RAM 420 may be used to implement one or more of network address translation tables present in system 400.

Secondary memory 430 may contain units such as hard drive 435 and removable storage drive 437. Secondary storage 430 may store the software instructions and data, which enable system 400 to provide several features in accordance with the present invention.

Some or all of the data and instructions may be provided on removable storage unit 440, and the data and instructions may be read and provided by removable storage drive 437 to processing unit 410. Floppy drive, magnetic tape drive, CD-ROM drive, DVD Drive, Flash memory, removable memory chip (PCMCIA Card, EPROM) are examples of such removable storage drive 437.

Processing unit 410 may contain one or more processors. Some of the processors can be general purpose processors which execute instructions provided from RAM 420. Some can be special purpose processors adapted for specific tasks (e.g., for memory/queue management). The special purpose processors may also be provided instructions from RAM 420. In general processing unit 410 reads sequences of instructions from various types of memory medium (including RAM 420, storage 430 and removable storage unit 440), and executes the instructions to provide various features of the present invention.

The description is continued with reference to an example implementation of server 130 substantially using integrated circuits.

9. Server

Figure 5 is a block diagram illustrating the details of an embodiment of server 140 as relevant to several aspects of the present invention. Server 130 is shown containing network interface 510, parser 520, SIP server block 530 and configuration block 540. Each block is described below in further detail.

Network interface 310 provides an electrical and protocol interfaces to receive and send internet protocol (IP) packets on network 111. Network interface 510 forwards received packets to parser 520. Similarly, packets received from configuration block 540 are transmitted on network 111. Network interface 310 may be implemented in a known way.

Parser 520 examines IP packets received from network interface 510 to determine the packets

which are to be sent to SIP server block 530. If a received packet relates to SIP, the packet is sent to SIP server block 530 for further processing.

SIP server block 530 processes the packets consistent with RFC 2543. In addition, SIP server block 530 examines the packet for the optional tag (briefly described in section above) in SIP registration messages to determine the old layer-3 address. SIP server block 530 then passes the old and new (also contained in the SIP registration message) IP addresses to configuration block 540.

Configuration block 540 configures (NAT entries and static route) router 120-B using the old and new IP addresses received from SIP server block 530. Protocols such as SNMP may be used for the configuration. In general, a configuration message in the form of one or more IP packets may be sent using network interface 510 during such configuration.

Thus, server 140 of Figure 5 can be used to configure router 120-B. It should be understood that when server 140 is integrated with router 130 into one unit/box, configuration block 540 may merely need to store a NAT entry in a NAT table (not shown) internal to the unit/box.

Server 170 may also be implemented similarly to configure the NAT entries in router 120-C. Once the two routers are configured, any sessions active prior to the move of mobile node 110-B may continue to operate without termination.

10. Conclusion

While various embodiments of the present invention have been described above, it should be

understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

Patent